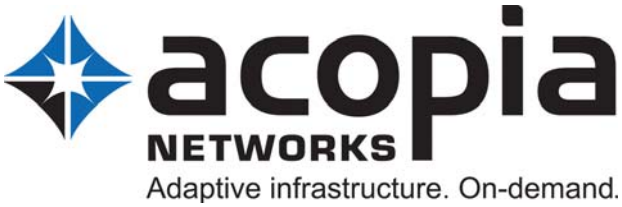

F5 Acopia ARX®
Software Release Notes
Release V2.07.001
Build 10312

F5 Acopia Networks®, Inc.

41 Wellman Street

Lowell, MA 01851

Part Number: 810-0024-00

Revision: V

Date: March, 2008

CONTENTS

Release Abstract

This document serves as the Release Notes for the Acopia FreedomFabric™ Network Operating System (NOS), version 2.07.001 (shortened to "2.7.1" in these notes). The FreedomFabric™ NOS runs on the Acopia Adaptive Resource Switch (ARX®).

Console Access

To access the Command-Line Interface (CLI) through the Console interface, use a terminal-emulation program (such as HyperTerminal) and a console cable provided by Acopia Networks. Set the terminal-emulation program to 9600-8-N-1, and connect the cable to the Console port on the ARX front panel.

Documentation

The following documentation is available in the installation kit:

- *Site Planning Guide* (Part number: 810-0027-00)
- *ARX®500 Hardware Installation Guide* (Part number: 810-0023-00)
- *ARX®1000 Hardware Installation Guide* (Part number: 810-0006-00)
- *ARX®6000 Hardware Installation Guide* (Part number: 810-0001-00)
- *GUI Quick Start: Network Setup* (Part number: 810-0042-00)
- *GUI Quick Start: NFS Storage* (Part number: 810-0049-00)
- *GUI Quick Start: CIFS Storage* (Part number: 810-0050-00)
- *CLI Network-Management Guide* (Part number: 810-0043-00)
- *CLI Storage-Management Guide* (Part number: 810-0044-00)
- *CLI Maintenance Guide* (Part number: 810-0045-00)
- *CLI Reference* (Part number: 810-0025-00)
- *SNMP Reference* (Part number: 810-0041-00)
- *Secure Agent Installation Guide* (Part number: 810-0013-00)
- *Log Catalog* (Part number: 810-0029-00)
- *Master Glossary* (Part number: 810-0011-00)
- *Master Index* (Part number: 810-0030-00)

You can access all of these manuals from the GUI: click on the **Documentation** link in the left-hand navigation panel.

New Features

Release 2.7.1 is a maintenance release for release 2.7.0. This release corrects anomalies in the 2.7.0 release. It does not offer new features.

Upgrade Paths

You can upgrade to 2.7.1 from any of the following releases:

- 2.4.3,
- 2.5.1,
- 2.5.2,
- 2.6.0, or
- 2.7.0.

New/Updated Certifications

There are no new client or file-server operating systems certified for use with the ARX.

For customers upgrading from Release 2.4.3, some certifications were added as of release 2.5.1. Acopia has certified the following platforms for use with the ARX:

- Microsoft Vista Business Edition clients.

  This is supported for Kerberos-only namespaces. The ARX does not support Microsoft Vista Home Edition.

- IBM Total Storage NAS 300G V2.0.1 file servers.

  This is supported for CIFS-only namespaces.

- Pillar Data Axiom V2.05 file servers.

  This is supported for CIFS-only namespaces and NFS-only namespaces.

- Acopia Secure Agent (ASA) support on Windows2003 SP2 platforms.

Firmware

This release does not supersede the firmware upgrades from release 2.5.1.

However, there is a firmware upgrade for installations that are upgrading from Release 2.4.3 (which pre-dated release 2.5.1). The firmware upgrade is for the ARX1000. After 2.7.1 is installed, you can use the firmware upgrade command to install the firmware. For detailed syntax and guidelines, refer to the Software Upgrade chapter in the *CLI Reference*.

2.7.0 Features

This section applies to installations that are upgrading from Release 2.6.0 or earlier.

Release 2.7.0 added three features, also included in this release:

- Persistent statistics for shadow-copy rules after failed runs of the rule,
- enhanced resiliency for shadow-copy runs, and
- improvements to the output for the show active-directory CLI command, to help with diagnosis of Active-Directory issues.

2.6.0 Features

This section applies to installations that are upgrading from Release 2.4.3 or 2.5.x.

Release 2.6.0 added two features, also included in this release:

- **Forest-to-Forest Trusts (Kerberos)** – Windows 2003 servers can now support trust relationships between AD forests. This allows clients in one of the AD forests to access a CIFS service in the other forest. In Release 2.6.0, the ARX can support these forest-to-forest trusts, allowing clients from a remote AD forest to authenticate with services in the local AD forest.
- **Shadow-Copy Support for CIFS Open Files** – By default, a shadow-copy rule blocks any writes from CIFS clients while it opens a file to copy it. No other client can write to the file or delete it in the middle of the copy operation. This prevents any corruptions in the shadow-copy of the file. However, if some other application is already holding the file open for writes before the shadow-copy rule opens it, the shadow-copy rule cannot open the file with read-only access. In Release 2.6.0, you can use the optional allow-shared-access command to successfully shadow copy any such open files.

2.5.1 Features

This section applies to installations that are upgrading from Release 2.4.3.

Release 2.5.1 offered these enhancements, also included in this release:

- **Forcing a Volume to "Take Ownership" of a Back-End Share** – A managed volume can now take ownership of a filer share that is evidently (but not actually) owned by another managed volume. This is an option for the CLI command, enable, for a share (in gbl-ns-vol-shr mode); refer to the CLI documentation for use cases and other details.
- **Flexible Naming for CIFS Services** – CIFS services now support NT domain names that do not match their AD domain names. Refer to the documentation for the CLI command, windows-domain (gbl-gs).
- **Performance Enhancements** – for removing shares from a managed volume.
- **Latency Statistics for Quorum Disks** – The CLI command, show redundancy quorum-disk, now shows latency statistics for the quorum disk. This is useful for anticipating redundancy-related issues.

- **Time Windows for Log Collection** – You can focus the collection of log messages by choosing an optional start time and end time. Refer to the documentation for the CLI command, collect logs.

- **Monitoring for CIFS-Client Activity** – A new CLI command, show cifs-service client-activity, shows details about a CIFS-client connection to the ARX, as well as the proxy connections from the ARX to the filers behind it.

2.5.0 Features

This section applies to installations that are upgrading from Release 2.4.3.

Release 2.5.0 offered the following new features, also included in this release:

- **Faster Imports into Managed Volumes** – In previous releases, managed volumes divided their imports into multiple scans, including a directory scan and a longer file scan. Volumes run the scans simultaneously in this new release. Additionally, a managed volume does not spend time protecting its metadata during the import; metadata protection is rarely necessary until the volume is running. Both options are configurable. Managed-volume imports are significantly faster with these options enabled, along with some other internal optimizations.

- **Support for CIFS subshares and their ACLs** – This release supports CIFS subshares and their share-level ACLs. A *subshare* is a CIFS share that is contained inside another share. The top-level share often has a different share-level Access Control List (ACL) than each of its subshares. By default, a volume always accesses directories through the top-level share, even when a client connects to a subshare on the front end. The volume connects to the top-level share, subjecting the client to the ACL there, then descends to the desired subdirectory. With this new feature, a CIFS service and CIFS volume can pass a client from a front-end subshare directly to the corresponding back-end subshare. The client therefore uses the subshare's ACL.

- **No Client Restrictions in Multi-Protocol (NFS and CIFS) Volumes** – Clients of a multi-protocol managed volume can now create directories with any name. In previous releases, volumes did not permit any names that resembled filer-generated names (FGNs, such as "myDir~1").

- **Ability to Migrate a Metadata Share** – There is a new option to migrate a managed volume's metadata from one dedicated share to another.

- **Firmware Upgrades** – You can use a new CLI command, **firmware upgrade**, to install any new firmware bundled with the latest software release. This command is documented in the software-upgrade chapter of the *CLI Reference*.

Automatic Upgrades for Multi-Protocol Volumes

A multi-protocol volume keeps additional metadata about its files and directories in the 2.5.0+ releases. The volume typically writes this information as it imports the volume's shares; this is not possible if you upgrade a previously-imported volume from Release 2.4.3. An automatic-upgrade process rewrites the metadata after the software upgrade.

The ARX runs an upgrade process in the background. It begins after the ARX boots with the 2.5.0+ software; if the ARX has a redundant peer, the upgrade begins after both peers are running the new release. The upgrade process runs on each multi-protocol volume, one at a time, to minimize CPU and bandwidth consumption. The process generates one report per volume, which you can access from the CLI with the

show reports command. Clients experience degraded performance in these volumes until the upgrades are completed.

After the upgrade, any directory with "~" in its name is accessible to NFS clients only. If your CIFS clients require access to one or more of these directories, run sync files for each of them. The sync files utility checks the volume's back-end filers for any FGNs that collide with the directory name(s), and it clears the NFS-only status if no collision is found.

For a Redundant Pair: Additional Reload after the Upgrade

In a redundant pair of ARXes, the above upgrade process only runs if both peers are running 2.5.0+ on startup. This is a safeguard to protect a software rollback: if the automatic-upgrade process changed all of the metadata to the 2.5.0+ format while the peer was running 2.4.3, a reversion to 2.4.3 would be impossible. Additionally, there are other irreversible 2.5.0+ features that only activate when both peers are running 2.5.*x*. Therefore, you must reboot the active peer after both peers are fully upgraded to 2.7.0. The *CLI Maintenance Guide* describes the full software-upgrade procedure. This causes a brief service outage while the processes fail over.

Problems Corrected

In This Release

- CIFS clients in a tree domain may be unable to authenticate to a VIP in another domain, even though the parents of the two domains have a valid two-way trust relationship. (27887)
- If a file-placement rule receives an excessive number of inline notifications (i.e., notifications of client changes), volume software may restart. (28770)
- If the same share name is used in multiple volumes, the **nsck … report** commands cannot always find the desired share. (28769)
- A **sync shares** operation can cause multiple ARX reboots if applied to a disabled volume with enabled shares. (29302)
- In rare cases, **nsck … report inconsistencies** fails to report some found directories (marked as "FD" in the inconsistencies report) until after a client renames or deletes them. (28708)
- When a CIFS client copies large directories into an ARX volume, the copy operation occasionally fails. (28735)
- In a presentation (or "direct") volume with an *attach point* to a managed volume, NFS clients hang when they attempt to create a hard link within the attached managed volume. (29106)
- In a presentation (or "direct") volume with an *attach point* to a managed volume, some CIFS directory listings get an incorrect directory path from the ARX volume. (29104)
- NFS3 UDP clients occasionally get incomplete directory listings from **ls**. (28668)
- If a directory in one CIFS share (for example, "longnameddirectory") has an alternate '8.3' name (e.g., "LONGNA~1") that collides with a directory in another CIFS share (with an actual name of "LONGNA~1"), the **sync files** command mistakenly assumes that all of the files in the first directory

("longnameddirectory") actually reside on the other share's directory (with the actual name of "LONGNA~1"). Clients cannot access the files unless someone runs another **sync files** operation. (28748)

- A **ron evict** command can result in a growing packet storm among the RON's remaining ARXes, possibly resulting in one or more ARX reboots. (28889)

- If a volume's CIFS client attempts to set the FILE_ATTRIBUTE_OFFLINE attribute (often used in stubs for filer-ILM applications), and one of the volume's filers cannot accept it, the volume may lock out all client access.
  Current Solution: Specific errors in the syslog indicate that this attribute was rejected by a particular filer, and F5/Acopia personnel have tools to work around the issue. Contact F5 Support if a CIFS volume blocks out client access and you see any TRANFS-…-INVOFFLINE or TRANFS-…-INVPARAM errors in the syslog. (28829)

- A **no share**, **remove-share**, or similar share-removal operation can be excessively slow in a large shadow volume, or in any managed volume with thousands of link files. (28494)

- A shadow volume may copy more files than were selected by the source fileset. (29021)

- If a file or directory in a source volume is renamed between shadow copies, and the new name is different from the old name in case only (e.g., "myfile.txt" becomes "MYFILE.txt"), the later shadow copy fails with a STATUS_OBJECT_NAME_COLLISION error. This only occurs in a shadow volume that supports CIFS. (28809)

- MTU path discovery is not negotiated properly. (28807)

- Configuring the same XIPs on both switches can prevent a VIP from starting. (28804)

- During DNS lookup, the domain name is translated to lowercase regardless of what is supplied by the user or configured on the switch, and the lookup fails. (28808)

In Release 2.7.0

- A memory fault in the Physical Address Extension (PAE) can lead to unpredictable issues, including an nsck destage that never completes, new VIPs that are unreachable, or even a chassis reboot. (27455)

- The **slot erase** command does not properly clean up the configuration database. (27756)

- Shadow-volume statistics restart from 0 (zero) after a shadow-copy run is interrupted by a failure (such as a reboot). (27046)

- The initial walk of a shadow-copy rule runs to completion after encountering a common error (such as a file opened by a CIFS client), then it restarts from the beginning. (27651)

- Volume software can fail (and write a core file) when a CIFS client has files open in two different volumes and the client's connection is lost. (27678)

- Certain client applications may crash unless the CIFS path-cache feature is disabled. (27013)

- NFS clients for a presentation volume experience occasional timeouts when accessing storage that is attached to a managed volume in the same VPU. (27934)

- NSM processors occasionally fail with a core file. (27987)
- Cannot press <Enter> for prompts through customer's Java SSL terminal. (28368)
- Metadata-Inconsistency reports do not report large directories correctly. (27229)
- An ARX with a large number of exports (more than 7200) sends an excessive number of nsmResourceThreshold traps. This resource trap is inappropriate for exports. (28442)
- The **delete capture** command causes a "Dynamic SQL Error" in the syslog. (27806)
- The ARX sends spurious vcifsSearchNotFound traps when CIFS clients search for filenames with invalid CIFS characters (such as ":"). (28210)
- A Windows **rmtshare** query fails on a valid ARX CIFS share. (28091)
- Pasting a running-config into the CLI generates errors within the cfg-channel commands. (28187)
- After a failover, many place rules report an inability to retrieve free-space information (shown in the **show policy** output). (28288)
- The output of the **collect diag** does not contain the output of the **show vpu** command. (28280)
- An ARX where all configuration was removed can experience reboots due to software faults. (27910)
- A place rule stops trying to migrate an open file after only five retries (two minutes and thirty seconds total); the timeout should be much longer. The solution was to increase the timeout to several hours. (28315)
- The Windows **rmtshare remark** command fails on a valid ARX CIFS share. (28381)
- The **policy pause** command and its GUI equivalent leave a share farm in an irreversible state, "Manually Paused." (28080)
- In a multi-protocol (CIFS and NFS) volume, file migrations to a UNIX Qtree caused UNIX permissions to change. This applies to Qtrees on NetApp 7.2.x filers. (28173)
- A place rule with duplicate inline notifications (migrations triggered by some client action) can reach a state where it perpetually restarts its volume scans. (27488)
- A managed volume fails with a core file if you remove one of its shares with the **force** option, then import another share while clients access the volume. (28233)
- The **show namespace** report shows a cryptic import error, "DTFS Operation has error status (-51)," when the back-end filer does not allow sufficient TCP sessions for the import to succeed. (28324)
- In environments where client and server clocks are unsynchronized, Kerberos authentications fail persistently for some applications. (25154)
- File migrations can stall indefinitely as they repeatedly attempt to get inaccessible directory locks. (27966)
- If a client modifies a file's CIFS SD through a presentation volume attached to a managed volume, a subsequent copy may result in a dual failure for both redundant ARXes. (27864)

- The policy engine continues running on the backup switch in a redundant pair, resulting in multiple "POLICY_PDP-0-3-POLICY_DB_OP_FAILED" messages in the syslog. (27875)

- Clients are unable to access ARX CIFS services through an SSL VPN. (27729)

- Kerberos authentications/logins are excessively slow (up to two minutes long) on an extremely busy ARX. (27983)

- NSM processors occasionally fail, writing a core file, when they have a CIFS connection to an EMC file server. (27739)

- The GUI is unresponsive when adding CIFS exports (28247)

- If you remove a volume and recreate it, the volume may stay in "Starting" state indefinitely. (27544)

- The **show active-directory status** command does not necessarily display active DCs that the ARX is using for Kerberos authentications. (28179)

- The domain-join operation, when performed in the GUI, leaves the administrator's password in clear text in the procdat log file. (28530)

- The policy engine can get into a state where it uses too much memory, eventually causing the ARX to reboot. (27804)

- A shadow-copy rule occasionally causes a policy-engine failure, resulting in one or more core files. (27338)

- The ARX cannot support 64 VIPs and 24 VPUs at the same time. (28487)

- If an NFS client renames one file so that it overwrites another, a deadlock can result in the NFS Lock Manager (NLM) software. (27956)

- The ARX drops some SNMP shareOnline traps. (28093)

- A back-end file name with a mixture of UTF-8 and 8859-1 characters can cause volume software to crash and produce a core file. (28097)

- The **collect diag** CLI command (and its GUI equivalent) includes system reports in the collected data; the diag option is supposed to omit reports, to keep the collected data at a reasonable size. (28331)

In Release 2.6.0

- Stale MAC addresses are not removed from a network processor after the processor returns from a failure. (27953)

- The ARX tends to reboot if its redundant peer fails and it has an unreliable quorum disk. (27137)

- In an unreliable network, NFS clients may cause dual failures in a redundant pair of NSM processors. This results in a failover to the backup ARX. (26100)

- MMC queries for CIFS-user sessions may incorrectly get an empty list of user sessions from an ARX-CIFS service. (27165)

- ARX volumes declare their back-end shares "offline" when write probes fail for 40 seconds. Multiple probe types should fail over a longer period for a share to be declared offline. (27400)

- A volume with **cifs path-cache** enabled rejects all requests for the "/" path (such as mount requests) from Mac Dave clients. (27404)

- If a managed volume's metadata filer is slow for an extended time, the volume may spontaneously re-import. (27178)

- If enough NFS clients hold open TCP connections to a VIP for an extended time, rpcinfo queries may fail and possibly cause a failover. (26829)

In Release 2.5.2

- When a thread fails during a managed-volume import, the other threads continue instead of aborting the import job. (27406)

- An NSM processor may fail if it is running nearly its maximum number of CIFS sessions, 16,000. (26496)

- The no capture session command can, in rare circumstances, lead to a system failure. (26844)

- An expect command can hang for an indefinite time. (26470)

- CIFS clients cannot access directories whose names begin with a ".", such as ".myFiles," while **cifs path-cache** is enabled. (27012)

- Some Excel-spreadsheet Macros fail when accessing the Excel files through the ARX. (26812, 26773)

In Release 2.5.1

- In installations with high-numbered IP addresses (140.x.x.x or higher) and many simultaneous connections, the ARX periodically loses contact with a back-end filer. (24423, 25956)

- A misleading redundancy error may appear when reloading the active peer. (25155)

- The SSH key changes after the software upgrade, requiring administrators to update their copy of the key. (26184, 26334)

- (Enhancement) Increase the maximum number of CIFS exports to 9,000. (26543)

- File or directory creates occasionally failed in CIFS managed volumes. (26624)

- CIFS clients are unable to access a file or directory with a German umlaut in its name. (26585)

- At a particular moment during startup of a Kerberos-and-NTLM service, it is possible for a Kerberos client to cause a CIFS-service failure. (26580)

- Alternate names are not supported in CIFS volumes. (24765)

- In the show namespace output for a namespace that is importing, the import status shows "N/A" for the first few seconds. (12377)

- A misleading error may appear when using the GUI to remove a volume. (25096)

- The Create Virtual Service wizard does not allow you to change the service's namespace after you have defined one export from it. (25107)

- You cannot enter an Organizational Unit (OU) with a name longer than 71 characters from the GUI. (25108)

In Release 2.5.0

- Report-prefix names cannot contain slash (/) characters. (22395)

- A misleading message appears when you use the CLI to remove the redundancy-link channel. (20326)

- The Disable button is ineffective in the GUI Namespace screen. (22202)

- If a volume's only metadata share is unavailable during import, you cannot nsck … destage the volume. (22360)

- After a failover, the show namespace command sometimes triggers a DB_TRANSACTION_CONFLICT error. (26234, 22922).

- If a simple-age fileset starts on the 29th-31st, shorter months confuse the scheduler. (20457)

- A file-placement rule configured by a script may hang in "Initializing" state. (24620)

- A shadow-copy rule does not duplicate changes in a directory's alternate data streams (ADS) if they change after the first run of the rule. (22598)

- When namespace-level metadata share goes offline, show namespace continues to display an "Online" metadata-share status in most of the volumes that use it. (21346)

- If the ARX reboots during a migration, the "pending" status of the migration is lost when the ARX comes back up. (21335)

Known Anomalies

This release contains the following known anomalies.

Problem	Description

Alarming error messages and logs appear during a firmware upgrade. (19365, 26774)	The following message appears in the CLI after you invoke and confirm a firmware upgrade:
	`ARXa6K# firmware upgrade all`
	`Confirmation of this command commences a firmware upgrade on the entire chassis. During the upgrade process, the chassis reboots automatically to complete the upgrade process.`
	`Proceed? [yes/no] yes`
	`ARXa6K#`
	`REDUNDANCY requires this switch to reboot.`
	`Reason: Paired NSM cores (0x530031, 0x530021) both failed.`
	`Switch reboot is prevented due to a firmware upgrade in progress. The system will retry every 10 seconds. Service failover is delayed until firmware upgrade is complete. Standby.`
	These messages are expected. The NSM processors (called "cores") are rebooting to load the new firmware. Additionally, the ARX sends SNMP traps (and, possibly E-mail alerts) to alert you to the rebooting processors.
	You can safely ignore all of these alerts until the firmware upgrade is complete. You can use the show firmware upgrade command to check the current status of the firmware upgrade.

Incorrect GUI prompt appears when you remove a Windows Mgmt. ACL. (26107)	If you go to Authentication -> Windows Mgmt. ACLs in the GUI, you have the option to remove any of the ACLs. An incorrect prompt appears in the confirmation pop-up; it warns about deleting an NTLM Auth. Server instead of the Windows Mgmt. ACL.
The show global-config output does not contain permit all any entries for windows-mgmt-auth groups. (26557)	When the global-config is replayed, members of a Windows-Management Authorization (WMA) group may lose all access through MMC. That is, a group that previously had full access would be left without any permissions to access the ARX volume(s) through MMC.

Problem	Description

The CLI show clock output does not always show the correct time after a time-zone change. (24526)	You can use the clock timezone CLI command to set the time zone of the ARX. On rare occasions, the output from the show clock command does not show the correct time after this change. For example:
	`ARXa500# clock set 14:43:00 01/11/2007`
	`ARXa500# show clock`
	`Local time: Thu Jan 11 14:43:02 2007 EST -0500 America New_York`
	`Universal time: Thu Jan 11 19:43:02 2007 UTC`
	`ARXa500# config`
	`ARXa500(cfg)# clock timezone America Denver`
	`ARXa500(cfg)# show clock`
	`Local time: Thu Jan 11 14:43:13 2007 EST -0500 America Denver`
	`Universal time: Thu Jan 11 19:43:13 2007 UTC`
	The time does not conform to the new time zone, though the correct new time zone (America Denver) does appear in the output.
	
	Workaround: Log out of the CLI and log back in.

During the hour of transition from daylight-savings time to standard time, the clock set CLI command incorrectly interprets times in some time zones. (24709)	Times are ambiguous in the hour when daylight-savings time reverts to standard time, once per year. Suppose the transition occurs at 3 AM on the day of the daylight-savings change: time passes from 3 to 4 AM in daylight-savings time, then the clock goes back to 3 AM for standard time, and then time passes from 3 to 4 AM again. In some time zones, if you reset the clock to a time between 3 and 4 AM, the clock set command may not interpret your time correctly. If this occurs, the ARX assumes that the transition to standard time has already occurred. This only occurs in time zones that are East of the Prime Meridian, with positive offsets from UTC. **Workaround**: Avoid the clock set command during the day and hour of transition.

Layer 3

Problem	Description
The ARX cannot send E-mail messages through the out-of-band (OOB) management interface. NTP, DNS, and RADIUS are also unsupported through the OOB interface. (24595)	All e-mail notifications from the ARX go out through an in-band (VLAN) management interface, configured with the interface vlan CLI command. At least one in-band-management interface must have a route to the E-mail server for E-mail notifications to function. The same applies to NTP, DNS, and RADIUS services. Use the cfg-mode ip route command (without the mgmt flag) to add a static IP route to the E-mail server(s), NTP server(s), and/or DNS server(s).

Redundancy

Problem	Description
Spurious errors appear in the syslog after an NSM failover. (25782)	NSM processors have redundant peers, even in an ARX that is not configured for overall redundancy. If an NSM processor fails, its peer processes packets for both. If **nsm recovery** is configured, the failed processor comes back online and waits to take over for the running processor. The failed processor may repeatedly put the following message in the syslog: `NAT rule TCP/`*`ip-address`*`:`*`port`*` for remote action `*`ip-address-2`*`:`*`port-2`*` type 3 not found.` This syslog message is spurious.

SNMP Traps and E-Mail Notifications

Problem	Description
The SCM may send redundant nvramECCError traps. (24580)	If two of these traps appear at roughly the same time, one from the SCM and one from the ASM, it is possible that the SCM trap is redundant. The correct number and disposition of ECC errors appears in the ARX's kernel log. Log into the CLI and use show logs kernel.log to view this log. ECC errors resemble the following: `Jan 10 03:10:41 sm-3-1 Uhhuh. NMI received. …`

Namespaces

Problem	Description
You must separately export a CIFS managed volume if you use it as a managed volume in a CIFS direct volume. (21231, 24359)	If a CIFS-managed volume is used as a managed volume in a CIFS-direct volume, its CIFS front-end service must export the managed volume separately. This is in addition to the export for the direct volume. (The same CIFS service must export both volumes.)

NSCK and Sync-Files

Problem	Description
NSCK reports do not identify "marked" multi-protocol directories where you should run a sync files operation. (23891)	Some multi-protocol (NFS and CIFS) directories are "marked" for special processing. These directories contain files and/or subdirectories one of these naming issues: • the name resembles a Filer-Generated Name (FGN, such as "myfile~1.txt"), or • the name produces an FGN on its back-end filers (such as "my:file.txt," or "MYFILE" in the same directory as "myfile"). If a directory is marked with one of these naming issues, the volume performs extra processing whenever a client tries to introduce an entry with the *other* naming issue. Depending on the outcome of the processing, the new client entry could become NFS-only (inaccessible to CIFS clients). Refer to the *CLI Maintenance Guide* for details. Clients can resolve these issues by accessing the volume through its VIP and renaming the directory's entries. However, the directory mark persists after all of its child entries have been correctly renamed; you use the sync files CLI command to remove the mark. The issue is that there are no reports that identify a directory as "marked" after its entries have been correctly renamed. **Workaround**: Use sync files to clear the directory mark immediately after renaming its entries.
Under rare circumstances, an nsck … rebuild on a shadow volume can make the volume stall in "importing" state. (18135)	If a shadow volume meets all of the following criteria when someone issues an nsck … rebuild, the shadow volume stays in "importing" state for a long time (perhaps hours), and is inaccessible to clients: • The shadow-copy-rule has its publish command set to group (the default), • The source volume contains millions of files, • The shadow-copy rule was near the end of a run, so most of the files were copied into the shadow volume's staging area (the hidden .acopia_shadow directory in the root of each share). The root of the problem is that the .acopia_shadow directories contain millions of files, and the nsck … rebuild must remove those directories at the beginning of its process. Clients cannot access the volume until all the filers are able to delete this directory. If this occurs, messages appear in the syslog that describe the problem.

Shadow Volumes

Problem	Description

A cryptic error appears in the shadow-copy report when the source volume is undergoing an nsck … rebuild. (20148)	If the source volume is being rebuilt during a shadow copy, the shadow copy fails appropriately. However, this error in the shadow-copy report is difficult to interpret:
	`target-switch: target-namespace:/target-vol/: /file`
	`% ERROR:  (23):`
	`  Read source file check data failed; File[/src-vol/file]`
	`Lookup of vfh`
	`failed; -17`